

CLAIMS

1 1. A system adapted to correct multiple storage device failures in a storage array using a
2 combination of multiple first parity groups and a single secondary parity group, the sys-
3 tem comprising:

4 a storage array having a plurality of concatenated sub-arrays, each sub-array in-
5 cluding a set of data storage devices and a first parity storage device, the array further
6 including a global secondary storage device associated with the storage array and holding
7 secondary parity values for the single secondary parity group, the secondary parity values
8 computed across the concatenation of the sub-arrays.

1 2. The system of Claim 1 wherein the first parity group is a row parity group, the first
2 parity storage device is a row parity storage device and wherein each row parity group is
3 associated with a sub-array of the storage array such that the array is composed of the
4 multiple row parity groups.

1 3. The system of Claim 2 wherein the secondary parity group is a diagonal parity group,
2 the secondary storage device is a diagonal parity storage device and wherein the secon-
3 dary parity values are diagonal parity values.

1 4. The system of Claim 3 further comprising:

2 a storage operating system configured to implement double failure protection en-
3 coding of the concatenated sub-arrays, wherein row parity values for each sub-array are
4 stored on each row parity storage device and diagonal parity values for the entire array
5 are stored on the global diagonal parity storage device; and

6 a processing element configured to execute the storage operating system to
7 thereby invoke storage access operations to and from the array in accordance with the
8 double failure protection encoding.

1 5. The system of Claim 4 wherein the double failure protection encoding is row-diagonal
2 parity encoding.

1 6. The system of Claim 4 wherein the double failure protection encoding is EVENODD
2 parity encoding.

1 7. The system of Claim 1 wherein each sub-array is organized as a concentrated parity
2 disk array.

1 8. The system of Claim 1 wherein each sub-array is organized as a distributed parity disk
2 array.

1 9. The system of Claim 1 wherein the storage devices are video tape, magnetic tape, opti-
2 cal, DVD, bubble memory, electronic random access memory or magnetic disk devices.

1 10. A method for correcting double failures in a storage array using a combination of a
2 single diagonal parity group and multiple row parity groups, the method comprising the
3 steps of:

4 organizing the storage array as a plurality of concatenated sub-arrays based on
5 double failure protection encoding, each sub-array including a set of data storage devices
6 and a row parity storage device, the storage array further including a global diagonal par-
7 ity storage device for holding diagonal parity;

8 computing the diagonal parity for the single diagonal parity group across the con-
9 catenated sub-arrays; and

10 correcting storage device failure within the array using the row parity storage de-
11 vice associated with each sub-array and the global diagonal parity storage device associ-
12 ated with the storage array.

1 11. The method of Claim 10 wherein the double failure protection encoding is row-
 2 diagonal parity encoding and wherein the step of correcting storage device failure com-
 3 prises the steps of:

4 determining whether the storage device failure is to a single storage device in a
 5 sub-array;

6 if the storage device failure is to a single storage device in the sub-array, recon-
 7 structing the failed storage device using local row parity associated with the sub-array;
 8 and

9 if the storage device failure is not to a single storage device in the sub-array, re-
 10 constructing the failed global diagonal parity storage device using all data and row parity
 11 storage devices of all sub-arrays of the array.

1 12. The method of Claim 11 wherein the step of correcting storage device failure further
 2 comprises the steps of:

3 if the storage device failure is not a single storage device failure, determining
 4 whether the storage device failure is a double failure within the sub-array;

5 if the storage device failure is not a double failure within the sub-array, determin-
 6 ing whether one of the failures includes the diagonal parity storage device; and

7 if one of the failures does not include the diagonal parity storage device, recon-
 8 structing the failed storage device in each sub-array using local row parity.

1 13. The method of Claim 12 wherein the step of correcting storage device failure further
 2 comprises the steps of:

3 if one of the failures includes the diagonal parity storage device, determining
 4 whether another of the failed storage devices includes a row parity storage device;

5 if the another of the failed storage devices includes the row parity storage device,
 6 reconstructing the row parity storage device from the data storage devices of the sub-
 7 array; and

8 reconstructing the diagonal parity storage device from all data and row parity
 9 storage devices of all sub-arrays of the array.

1 14. The method of Claim 13 wherein the step of correcting storage device failure further
2 comprises the steps of:

3 if the another of the failed storage devices does not include the row parity storage
4 device, reconstructing the data storage device using local row parity associated with the
5 sub-array; and

6 reconstructing the diagonal parity storage device from all data and row parity
7 storage devices of all sub-arrays of the array.

1 15. The method of Claim 14 wherein the step of correcting storage device failures further
2 comprises the step of, if the storage device failure is a double failure within the sub-array,
3 recovering two failed storage devices within the sub-array using a row-diagonal recon-
4 struction process.

1 16. The method of Claim 15 wherein the step of recovering comprises the steps of:
2 using the diagonal parity storage device to recover at least one data block from a
3 first of the failed storage devices of the sub-array; and
4 once the data block is recovered, using row parity within the sub-array to recover
5 a corresponding block in a second of the failed storage devices.

1 17. The method of Claim 10 further comprising the step of organizing each sub-array as
2 a concentrated parity disk array.

1 18. The method of Claim 10 further comprising the step of organizing each sub-array as
2 a distributed parity disk array.

1 19. The method of Claim 10 wherein the double failure protection encoding is
2 EVENODD parity encoding and wherein the step of correcting storage device failure
3 comprises the steps of:
4 determining whether the storage device failure is to a single storage device in a
5 sub-array;

6 if the storage device failure is to a single storage device in the sub-array, recon-
7 structing the failed storage device using local row parity associated with the sub-array;
8 and

9 if the storage device failure is not to a single storage device in the sub-array, re-
10 constructing the failed global diagonal parity storage device using all data storage devices
11 of all sub-arrays of the array.

1 20. The method of Claim 19 wherein the step of correcting storage device failure further
2 comprises the steps of:

3 if the storage device failure is not a single storage device failure, determining
4 whether the storage device failure is a double failure within the sub-array;

5 if the storage device failure is not a double failure within the sub-array, determin-
6 ing whether one of the failures includes the diagonal parity storage device; and

7 if one of the failures does not include the diagonal parity storage device, recon-
8 structing the failed storage device in each sub-array using local row parity.

1 21. The method of Claim 20 wherein the step of correcting storage device failure further
2 comprises the steps of:

3 if one of the failures includes the diagonal parity storage device, determining
4 whether another of the failed storage devices includes a row parity storage device;

5 if the another of the failed storage devices includes the row parity storage device,
6 reconstructing the row parity storage device from the data storage devices of the sub-
7 array; and

8 reconstructing the diagonal parity storage device from all of the data storage de-
9 vices of the array.

1 22. The method of Claim 21 wherein the step of correcting storage device failure further
2 comprises the steps of:

3 if the another of the failed storage devices does not include the row parity storage
 4 device, reconstructing the data storage device using local row parity associated with the
 5 sub-array; and

6 reconstructing the diagonal parity storage device from all data storage devices of
 7 the array.

1 23. The method of Claim 22 wherein the step of correcting storage device failures further
 2 comprises the step of, if the storage device failure is a double failure within the sub-array,
 3 recovering two failed storage devices within the sub-array using an EVENODD recon-
 4 struction process.

1 24. Apparatus for correcting double failures in a storage array using a combination of a
 2 single diagonal parity group and multiple row parity groups, the apparatus comprising:

3 means for organizing the storage array as a plurality of concatenated sub-arrays
 4 based on double failure protection encoding, each sub-array including a set of data stor-
 5 age devices and a row parity storage device, the storage array further including a global
 6 diagonal parity storage device for holding diagonal parity;

7 means for computing the diagonal parity for the single diagonal parity group
 8 across the concatenated sub-arrays; and

9 means for correcting storage device failure within the array using the row parity
 10 storage device associated with each sub-array and the global diagonal parity storage de-
 11 vice associated with the storage array.

1 25. A computer readable medium containing executable program instructions for correct-
 2 ing double failures in a storage array using a combination of a single diagonal parity
 3 group and multiple row parity groups, the executable program instructions comprising
 4 program instructions for:

5 organizing the storage array as a plurality of concatenated sub-arrays based on
 6 double failure protection encoding, each sub-array including a set of data storage devices

7 and a row parity storage device, the storage array further including a global diagonal par-
8 ity storage device for holding diagonal parity;
9 computing the diagonal parity for the single diagonal parity group across the con-
10 catenated sub-arrays;
11 correcting storage device failure within the array using the row parity storage de-
12 vice associated with each sub-array and the global diagonal parity storage device associ-
13 ated with the storage array.

1 26. A system adapted to correct multiple storage element failures in an array using a
2 combination of multiple first failure recovery groups and a single secondary failure re-
3 covery group, the system comprising:
4 a storage array having a plurality of concatenated sub-arrays, each sub-array in-
5 cluding a set of data storage elements and a first failure recovery storage element storing
6 first values used to correct a single failure within the sub-array, the array further includ-
7 ing a global failure recovery storage element associated with the storage array and hold-
8 ing secondary values for the single secondary failure recovery group, the secondary val-
9 ues computed across the concatenation of the sub-arrays.

1 27. The system of Claim 26 wherein the storage elements are packets and wherein the
2 failure recovery is parity.